Towards Creating Better Interactive Agents: Leveraging Both Implicit and Explicit Human Feedback

Doctoral Consortium

Kate Candon
Yale University
New Haven, CT, USA
kate.candon@yale.edu

ABSTRACT

My work aims to create interactive agents that are more effectively able to help people. The way in which people want to be helped can vary based on a number of factors, such as person, task, or time. Thus, an important capability of interactive agents is to be able to tailor their behavior based on a person's preferences throughout an interaction. Typically, interactive agents can learn a person's preferences from explicit feedback, such as evaluative (good versus bad) feedback, corrections, or demonstrations. However, there are downsides to relying only on explicit feedback. Therefore, it would be advantageous if interactive agents could also adapt to a person's preferences based on feedback provided implicitly. Implicit human feedback can include information such as eye gaze, facial reactions, or a person's own choice of actions in a task. This line of research investigates reasoning about both implicit and explicit human feedback together during an interaction. For example, we propose reasoning about implicit human feedback in order to proactively solicit explicit feedback. This could allow an interactive agent to proactively tailor its behavior to the preferences of the person with whom they are interacting.

KEYWORDS

human-agent interaction; implicit feedback; explicit feedback

ACM Reference Format:
Kate Candon. 2023. Towards Creating Better Interactive Agents: Leveraging Both Implicit and Explicit Human Feedback: Doctoral Consortium. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

This work aims to create interactive agents that are effectively able to help people in a personalized manner during continuous collaborations. Imagine you are working with a robot to prepare dinner. You typically like to keep your workstation as clean as possible, so you prefer to only gather ingredients when you need them and to clean up as soon as you're done with an item. Otherwise, you do not have strong preferences about how to split tasks with the robot. During the interaction, you might frown when the robot fetches an ingredient you do not yet want in the workspace or subconsciously signal when you are ready for a different ingredient by directing your gaze towards it. While you want the robot to adjust to your preferences, it would be both annoying and distracting if the robot asked for feedback after every single action (or inaction). How could the robot ensure it collaborates with you according to your preferences?

Figure 1: Individual playing Space Invaders with a Nao robot.

A trivial approach would be to assume all people have the same preferences when receiving help from an agent. Then, assistive actions could be pre-determined based on individual tasks. We could assign rewards based on the set up of an individual task, and the robot could select assistive actions based on which action would increase the reward of the person they are trying to assist.

However, in preliminary work [4], we found that even when participants were given the same instructions, they had different interpretations of and preferences between assistive behaviors. In an exploratory study, we had participants play a video game alongside an interactive agent (or *co-player*) displaying different types of assistive behaviors, and we studied the factors influencing the perceived helpfulness in the interactions. Specifically, we used a custom two-player version of Space Invaders (as seen on the computer screen in Figure 1). We chose Space Invaders because it requires continuous and fast-paced decision-making and action from participants. Also, the participant is continuously engaged in their own portion of the task while also being asked to provide feedback.

We found that even in a well-structured domain such as Space Invaders, the way in which participants interpreted assistive actions from the co-player was nuanced. Additionally, we found that helpfulness was more correlated with emotional perceptions (such as how annoying the participant viewed the co-player) than with game objectives (how many points the co-player scored for the player). These findings are in line with recent work in Artificial

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

Intelligence (AI) suggesting that the most accurate model is not necessarily the best in human-AI teams [2]. It is challenging to know for what an agent should be optimizing.

All together, these results suggest that the notion of assistance in nuanced and personal, so it is important for an agent to be able to adapt its behaviors when trying to help someone.

2 LEVERAGING NON-VERBAL BEHAVIORS

Due to the nuanced nature of preferences surrounding assistance, we believe it is important to leverage all available information during an interaction. Typically, agents learn to adapt to individual preferences via explicit human feedback [7]. Common types of feedback in include evaluative feedback [9], demonstrations [12], corrections [1], and comparisons [11]. However, different problems can arise if an agent relies only on explicit feedback. For example, humans tend to provide less explicit feedback as an interaction progresses [10] or may cease providing feedback once they are satisfied with an agent's performance [8]. Additionally, explicit feedback places additional burden on the human in the interaction and can interrupt the flow of a task.

Because of these limitations of explicit feedback, it would be advantageous if agents could use other implicit information to complement necessary explicit feedback. Implicit human feedback can include information such as eye gaze, facial reactions, or a person's own choice of actions in a task. An advantage of implicit feedback is that it is provided naturally during interactions.

We investigated the usefulness of naturally provided human nonverbal signals in understanding user preferences [5]. We used the images recorded via a participant's webcams while they played Space Invaders to see if we could build models that would predict which co-player behavior a participant preferred. We found that even without explicitly directing participants to be expressive, incorporating "free" nonverbal reactions improved our ability to predict their preferences between agent behaviors.

Additionally, our results suggest that considering additional context is important when trying to interpret nonverbal human behavior. Specifically, we found that incorporating demographic and personality information about the participant and/or game statistics, further improved our ability to predict which co-player behavior the participant preferred. This finding aligns with research in social psychology emphasizing the importance of context when trying to recognize emotions [3] and emphasizes that there are challenges to relying on implicit feedback alone.

3 CONSIDERATIONS FOR EXPLICIT FEEDBACK

While there is promise in leveraging implicit feedback to tailor agent behaviors, it it still important to understand how to incorporate explicit feedback into human-agent interactions. In another line of research, we explored the effect of reminders for a human to provide feedback about a robot's behavior during an interaction [6]. Specifically, we investigated the influence of how and when a robot reminded the participant to provide feedback. We found that when the robot framed a reminder as helping the team improve, participants felt more positively about the robot and about the process of providing feedback than when the robot framed a reminder as helping the robot improve its individual performance. Additionally, by reminding a participant to provide feedback before a change in behavior, a robot could influence participants to provide feedback more quickly and more frequently compared to when the reminder was after a change in behavior.

4 FUTURE DIRECTIONS

Motivated by our preliminary findings, the overarching goal of this line of work is to create interactive agents that are more effectively able to help people by tailoring their behavior to individual preferences during an interaction. We propose to achieve this by leveraging nonverbal cues to intelligently solicit explicit feedback.

First, we need to create models that can perceive changes in nonverbal behavior and reason about what those nonverbal behaviors may signal about a person's preferences. One approach is to build a model that can predict when an individual is about to provide positive or negative explicit feedback. This would enable us to predict a reward during an interaction, and thus we could lessen the requirement for the participant to explicitly provide feedback. Another approach is to build a model that could detect anomalies or changes in nonverbal behavior, which could signal that the user may have opinions about the actions the agent is taking.

Second, there are remaining open questions about how to best solicit and react to explicit feedback during interactions. For example, we could explore if the way in which an interactive agent responds to explicit feedback affects the future feedback a participant provides. We are also interested in exploring not only the frequency and timing of explicit feedback, but the quality of explicitly provided feedback in different scenarios.

Finally, combining a better understanding of both implicit and explicit human feedback, we can explore how an agent can most efficiently query a user for explicit feedback during a task. For example, if an agent perceives a nonverbal reaction that it suspects indicates a negative reward, the agent could either confirm its most recent action was in fact not desired, ask what it should have done instead, or ask why what it did was wrong. Alternatively, if an agent tries different actions, but the inferred reward remains low or there is no perceived change in nonverbal behavior, it may be a good time to ask the user for input.

5 CONCLUSION

We hope that our work towards a better understanding of implicit and explicit human feedback will encourage the community to think about how to leverage both types of feedback during interactions for proactive personalization. We believe that by building systems that reason about both types of feedback together, interactive agents will be able to more naturally adapt to individual preferences.

ACKNOWLEDGMENTS

Thank you to M. Vázquez for her consistent guidance and support as an advisor. Thanks to Z. Hsu, Y. Kim, J. Chen, H. Zhou, S. Gillet, and N. Tsoi for their collaboration on preliminary results discussed herein. This work was partially supported by the National Science Foundation (NSF) under Grant No. IIS-2106690. Any opinions, findings, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

[1] Andrea Bajcsy, Dylan P. Losey, Marcia K. O'Malley, and Anca D. Dragan. 2017. Learning Robot Objectives from Physical Human Interaction. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*. PMLR, 217–226.

[2] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. 2021. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11405–11414.

[3] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. 2019. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20, 1 (2019), 1–68.

[4] Kate Candon, Zoe Hsu, Yoony Kim, Jesse Chen, Nathan Tsoi, and Marynel Vázquez. 2022. Perceptions of the Helpfulness of Unexpected Agent Assistance. In *Proceedings of the 10th International Conference on Human-Agent Interaction* (Christchurch, New Zealand) *(HAI '22)*. Association for Computing Machinery, New York, NY, USA, 41–50. https://doi.org/10.1145/3527188.3561915

[5] Kate Candon, Zoe Hsu, Yoony Kim, Jesse Chen, Nathan Tsoi, and Marynel Vázquez. 2023. Nonverbal Human Signals Can Help Autonomous Agents Infer Human Preferences for Their Behavior. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*. IFAAMAS.

[6] Kate Candon, Helen Zhou, Sarah Gillet, and Marynel Vázquez. 2023. Verbally Soliciting Human Feedback in Continuous Human-Robot Collaboration: Effects of the Framing and Timing of Reminders. In *Proceedings of the 18th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.

[7] Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid Simmons, Aaron Steinfeld, and Tesca Fitzgerald. 2021. Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning. In *International Joint Conference on Artificial Intelligence*.

[8] Charles Isbell and Christian Shelton. 2001. Cobot: A Social Reinforcement Learning Agent. In *Advances in Neural Information Processing Systems*, Vol. 14. MIT Press. https://proceedings.neurips.cc/paper/2001/file/92bbd31f8e0e43a7da8a6295b251725f-Paper.pdf

[9] W. Bradley Knox and Peter Stone. 2008. TAMER: Training an Agent Manually via Evaluative Reinforcement. In *2008 7th IEEE International Conference on Development and Learning*. IEEE, 292–297.

[10] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. 2013. Using Informative Behavior to Increase Engagement in the Tamer Framework. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '13)*. IFAAMAS, 909–916.

[11] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. 2017. Active Preference-Based Learning of Reward Functions. In *Proceedings of Robotics: Science and Systems (RSS)*.

[12] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. 2022. MIND MELD: Personalized Meta-Learning for Robot-Centric Imitation Learning. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 157–165.